

Accelerating the Fourier split operator method via graphics processing units

Heiko Bauke* and Christoph H. Keitel†

Max-Planck-Institut für Kernphysik, Saupfercheckweg 1, 69117 Heidelberg, Germany

Current generations of graphics processing units have turned into highly parallel devices with general computing capabilities. Thus, graphics processing units may be utilized, for example, to solve time dependent partial differential equations by the Fourier split operator method. In this contribution, we demonstrate that graphics processing units are capable to calculate fast Fourier transforms much more efficiently than traditional central processing units. Thus, graphics processing units render efficient implementations of the Fourier split operator method possible. Performance gains of more than an order of magnitude as compared to implementations for traditional central processing units are reached in the solution of the time dependent Schrödinger equation and the time dependent Dirac equation.

PACS numbers: 02.70.-c, 03.65.Pm

I. Introduction

For several decades, users and developers of high performance computing applications could trust on Moore's law. The number of transistors that can be placed inexpensively on an integrated circuit has doubled approximately every two years. This exponential growth provided the basis for an ever-increasing computing power of modern central processing units (CPUs). Thus, high performance computing applications became faster by just trusting on Moore's law and waiting for faster CPUs. Thanks to Moore's law, today's desktop computers supply the computational power of the last decade's super-computers.

Transistor counts still continue to grow exponentially. Hardware manufacturers, however, no longer make single computing units more complex and faster. Clock rates and execution optimizations have reached their limits. Now hardware manufacturers use the plethora of transistors to put many computing units on a single chip manufacturing highly parallel computing architectures. There are two major directions of development, multicore architectures with a few (typically two to a few dozen in the near future) general purpose computing units and dedicated accelerators with several hundred or more computing units with reduced capabilities each as compared to traditional CPUs [1, 2]. These accelerators may be implemented in field-programmable gate arrays or may be custom built systems as the Cell Broadband Engine Architecture or the ClearSpeed™ CSX700 Processor, for example. Graphics processing units (GPUs) with general purpose computing capabilities brought accelerated computing to the mass market. Employing GPUs to perform computations that are traditionally handled by CPUs is commonly referred to as GPU computing [3].

Parallel architectures are not new in the high performance computing community. In the past, many high performance computing applications employed various kinds of parallel computing architectures which had been the high-end com-

puting systems of their time. Today, however, even low-end consumer computers are parallel systems. The new emerging ubiquitous parallel architectures indicate a silent paradigm change in computing or as Herb Sutter [4] pointed out "The free lunch is over." In order to exploit the power of modern hardware architectures it is required to write parallel applications. The term "computing" becomes synonymous to "parallel computing."

The purpose of this contribution is to evaluate the Fermi GPU computing architecture and to demonstrate how an implementation of the Fourier split operator method on GPUs may boost the performance of the numerical propagation of time dependent partial differential equations. The remainder of this paper is organized as follows. In section II we give a generic description of the Fourier split operator method and show how to apply it to the time dependent Schrödinger equation and the time dependent Dirac equation. Section III gives a very short introduction to the CUDA architecture and describes our GPU implementation of the Fourier split operator method while section IV presents performance comparisons. In section V we illustrate some applications of the Fourier split operator method.

II. The Fourier split operator method

Let us consider the Cauchy type initial value problem

$$\frac{\partial w(\mathbf{x}, t)}{\partial t} = \hat{A}(t)w(\mathbf{x}, t) \quad (1a)$$

with some possibly time dependent differential operator $\hat{A}(t)$ acting on the coordinates $\mathbf{x} = (x_1, x_2, \dots, x_n)$ and with the initial condition

$$w(\mathbf{x}, 0) = w_0(\mathbf{x}). \quad (1b)$$

The function $w(\mathbf{x}, t)$ depends on \mathbf{x} and the time coordinate t and may be real or complex and possibly vector valued. The Cauchy problem (1) includes, for example, the diffusion equation, the Fock-Planck equation [5], the current-free Maxwell

* bauke@mpi-hd.mpg.de

† keitel@mpi-hd.mpg.de

equations [6], the Schrödinger equation, the Pauli equation, the Dirac equation [7], as well as the Klein-Gordon equation in the Feshbach-Villars representation [8] or the Black-Scholes options pricing equation.

Using Dyson's time ordering operator \hat{T} , the formal solution of (1) is given by

$$w(\mathbf{x}, t) = \hat{U}(t, 0)w(\mathbf{x}, 0) \quad (2)$$

with the time-evolution operator [5, 9, 10]

$$\hat{U}(t_2, t_1) = \hat{T} \exp \left(\int_{t_1}^{t_2} \hat{A}(t') dt' \right) \quad (3)$$

that effects the function's evolution from time t_1 to time t_2 . For some highly symmetric operators $\hat{A}(t)$ in (1), analytic expressions of the time-evolution operator (3) can be calculated. However, investigations of many relevant problems require numerical methods, for example, the Fourier split operator method, which we will shortly describe in the following paragraphs.

II.1. General outline of the Fourier split operator method

Fleck et al. [11] introduced the Fourier split operator method as an explicit time stepping scheme for the solution of the time dependent scalar Maxwell wave equation in Fresnel approximation. The method is applicable to Cauchy problems (1) with a differential operator $\hat{A}(t)$ that has the property that it may be separated into a sum of two operators

$$\hat{A}(t) = \hat{A}_1(t) + \hat{A}_2(t) \quad (4)$$

such that $\hat{A}_1(t)$ can be easily diagonalized in real space while $\hat{A}_2(t)$ can be easily diagonalized in Fourier space; a requirement that is fulfilled by many partial differential equations of relevance. Thus, shortly after Feit et al. [12] solved the Schrödinger equation with a time independent Hamiltonian numerically by the Fourier split operator method, it became a standard tool for the propagation of quantum mechanical wave equations. Later, the method had been generalized to the Schrödinger equation with a time dependent Hamiltonian [13] and it was applied to other equations, for example, the Dirac equation [14–17], the time dependent Gross-Pitaevskii equation [18], the non-linear Schrödinger equation [19], and the time dependent Maxwell equations for electromagnetic waves in random dielectric media [6].

The central idea of the Fourier split operator method is to approximate the operator (3) by a product of operators that are diagonal either in real space or in momentum space. Let $\hat{O}(t)$ denote some possibly time dependent operator and define the operator

$$\hat{U}_{\hat{O}}(t_2, t_1, \delta) = \exp \left(\delta \int_{t_1}^{t_2} \hat{O}(t') dt' \right), \quad (5)$$

that depends on the times t_1 and t_2 and the auxiliary parameter δ and let $\hat{\hat{U}}_{\hat{O}}(t_2, t_1, \delta)$ denote the operator (5) in Fourier

space. Expanding $\hat{U}(t + \tau, t)$ to the third order in τ and provided that the operator $\hat{A}(t)$ has the form (4), the time-evolution operator (3) can be factorized [9] into

$$\begin{aligned} \hat{U}(t + \tau, t) &= \exp \left(\int_t^{t+\tau} \hat{A}(t') dt' \right) + O(\tau^3) = \\ &\hat{U}_{\hat{A}_1} \left(t + \tau, t, \frac{1}{2} \right) \hat{U}_{\hat{A}_2} (t + \tau, t, 1) \hat{U}_{\hat{A}_1} \left(t + \tau, t, \frac{1}{2} \right) + O(\tau^3). \end{aligned} \quad (6)$$

Neglecting terms of order $O(\tau^3)$, equation (6) gives an explicit second order accurate time-stepping scheme for the propagation of the function $w(\mathbf{x}, t)$

$$\begin{aligned} w(\mathbf{x}, t + \tau) &= \\ \hat{U}_{\hat{A}_1} \left(t + \tau, t, \frac{1}{2} \right) \hat{U}_{\hat{A}_2} (t + \tau, t, 1) \hat{U}_{\hat{A}_1} \left(t + \tau, t, \frac{1}{2} \right) w(\mathbf{x}, t) &+ O(\tau^3). \end{aligned} \quad (7)$$

This scheme translates the difficulty of calculating the action of operator (3) to the task of calculating the action of (5) for $\hat{O} \equiv \hat{A}_1$ and $\hat{O} \equiv \hat{A}_2$, respectively. Strang [20] utilized the splitting (7) to calculate the action of (5) in real space by a finite differences scheme. For many Cauchy problems (1), however, one can find a splitting (4) such that the operator $\hat{U}_{\hat{A}_1}(t + \tau, t, \delta)$ is diagonal in real space and $\hat{U}_{\hat{A}_2}(t + \tau, t, \delta)$ is diagonal in Fourier space. Thus, the calculation of these operators becomes feasible in the appropriate space and (7) is calculated via

$$\begin{aligned} w(\mathbf{x}, t + \tau) &= \\ \hat{U}_{\hat{A}_1} \left(t + \tau, t, \frac{1}{2} \right) \mathcal{F}^{-1} \left\{ \hat{U}_{\hat{A}_2} (t + \tau, t, 1) \mathcal{F} \left\{ \hat{U}_{\hat{A}_1} \left(t + \tau, t, \frac{1}{2} \right) w(\mathbf{x}, t) \right\} \right\} &+ O(\tau^3). \end{aligned} \quad (8)$$

The expression $\mathcal{F}\{\cdot\}$ in (8) denotes the Fourier transform of the argument and $\mathcal{F}^{-1}\{\cdot\}$ the inverse Fourier transform.

In a computer implementation of the Fourier split operator method, the function $w(\mathbf{x}, t)$ is discretized on a rectangular regular lattice of N points and the continuous Fourier transform is approximated by a discrete Fourier transform. The computational complexity of propagating the function $w(\mathbf{x}, t)$ from time t to time $t + \tau$ is dominated by the transformation into Fourier space and back into real space. If these transforms are accomplished by the fast Fourier transform the computation of an elementary step of the Fourier split operator method takes $O(N \log N)$ operations.

II.2. Fourier split operator method for the Schrödinger equation

The Schrödinger equation is an equation of motion for a complex-valued scalar wave function $\Psi(\mathbf{x}, t)$ evolving in a d -dimensional space \mathbf{x} and in time t . In its most general form, it describes a non-relativistic spin zero particle of mass m and charge q in the electromagnetic potentials $\phi(\mathbf{x}, t)$ and $\mathbf{A}(\mathbf{x}, t)$.

The Schrödinger equation is invariant under Galilean transformation and reads

$$i\hbar \frac{\partial \Psi(\mathbf{x}, t)}{\partial t} = \left(\frac{1}{2m} (-\hbar i \nabla - q\mathbf{A}(\mathbf{x}, t))^2 + q\phi(\mathbf{x}, t) \right) \Psi(\mathbf{x}, t). \quad (9)$$

In order to apply the Fourier split operator method for propagating $\Psi(\mathbf{x}, t)$ under the effect of (9), we have to restrict the vector potential to homogeneous fields $\mathbf{A}(\mathbf{x}, t) = \mathbf{A}(t)$; an approximation that is known as the dipole approximation. Splitting the Hamiltonian (9) in dipole approximation into a potential energy term and a kinetic energy term

$$i\hbar \hat{A}_1 = q\phi(\mathbf{x}, t), \quad (10a)$$

$$i\hbar \hat{A}_2 = \frac{1}{2m} (-\hbar i \nabla - q\mathbf{A}(t))^2 \quad (10b)$$

separates the spatial dependent parts from spatial derivatives, which makes the operator $\hat{U}_{\hat{A}_1}(t + \tau, t, \delta)$ diagonal in real space and $\hat{U}_{\hat{A}_2}(t + \tau, t, \delta)$ diagonal in momentum space, respectively. The action of $\hat{U}_{\hat{A}_1}(t + \tau, t, \delta)$ on a real space wave function is given by

$$\hat{U}_{\hat{A}_1}(t + \tau, t, \delta) \Psi(\mathbf{x}, t) = \exp\left(-\delta \frac{i}{\hbar} \int_t^{t+\tau} q\phi(\mathbf{x}, t') dt'\right) \Psi(\mathbf{x}, t) \quad (11)$$

and the action of the Fourier space operator $\hat{U}_{\hat{A}_2}(t + \tau, t, \delta)$ to a wave function in Fourier space

$$\tilde{\Psi}(\mathbf{p}, t) = \mathcal{F}\{\Psi(\mathbf{x}, t)\} = \frac{1}{(2\pi\hbar^2)^{d/2}} \int \Psi(\mathbf{x}, t) \exp(-i\mathbf{p} \cdot \mathbf{x}/\hbar) d^d x \quad (12)$$

reads

$$\hat{U}_{\hat{A}_2}(t + \tau, t, \delta) \tilde{\Psi}(\mathbf{p}, t) = \exp\left(-\delta \frac{i}{\hbar} \int_t^{t+\tau} \frac{1}{2m} (\mathbf{p} - q\mathbf{A}(t'))^2 dt'\right) \tilde{\Psi}(\mathbf{p}, t). \quad (13)$$

Note that it is crucial for the application of the Fourier split operator method that the vector potential $\mathbf{A}(t)$ does not depend on the spatial coordinate \mathbf{x} . The expansion of the Hamiltonian of the Schrödinger equation (9) for a particle in an arbitrary vector potential $\mathbf{A}(\mathbf{x}, t)$ contains the term $(iq\hbar/m)\mathbf{A}(\mathbf{x}, t) \cdot \nabla$, that is spatially dependent and contains spatial derivatives, too, coupling momentum and coordinate space. Coupling between momentum and coordinate space is absent for vector potentials in dipole approximation but also, for example, for the vector potential of a linearly polarized plane wave with $\mathbf{A}(\mathbf{x}, t) = (A_1(x_3, t), 0, 0)$ [17].

II.3. Fourier split operator method for the Dirac equation

In contrast to the Schrödinger equation, the Dirac equation [7] describes a relativistic spin half particle and it is invariant under Lorentz transformation. A Dirac wave function $\Psi(\mathbf{x}, t)$ is a

two component (for one-dimensional systems) or four component (for two- or three-dimensional systems) complex-valued vector function. The Dirac equation for a particle of mass m and charge q moving in the electromagnetic potentials $\phi(\mathbf{x}, t)$ and $\mathbf{A}(\mathbf{x}, t)$ is given by

$$i\hbar \frac{\partial \Psi(\mathbf{x}, t)}{\partial t} = \left(c \sum_{i=1}^d \alpha_i \left(-i\hbar \frac{\partial}{\partial r_i} - qA_i(\mathbf{x}, t) \right) + q\phi(\mathbf{x}, t) + mc^2 \beta \right) \Psi(\mathbf{x}, t) \quad (14)$$

with the matrices α_i and β and the speed of light c in vacuum and $A_i(\mathbf{x}, t)$ denoting the i th component of the vector potential $\mathbf{A}(\mathbf{x}, t)$. The matrices α_i, β obey the algebra

$$\alpha_i^2 = \beta^2 = 1, \quad \alpha_i \alpha_k + \alpha_k \alpha_i = 2\delta_{i,k}, \quad \alpha_i \beta + \beta \alpha_i = 0. \quad (15)$$

This algebra determines the matrices α_i and β only up to unitary transforms. In numerical applications, we adopted the so-called Dirac representation with

$$\alpha_1 = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \quad \alpha_2 = \begin{pmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \\ 0 & -i & 0 & 0 \\ i & 0 & 0 & 0 \end{pmatrix}, \quad (16)$$

$$\alpha_3 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}, \quad \beta = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}.$$

For one-dimensional systems, the Dirac representation reduces to

$$\alpha_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \beta = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \quad (17)$$

In order to apply the Fourier split operator method to the Dirac equation (14), we split the Hamiltonian into an interaction part and a free particle part

$$i\hbar \hat{A}_1 = c \sum_{i=1}^d \alpha_i (-qA_i(\mathbf{x}, t)) + q\phi(\mathbf{x}, t), \quad (18a)$$

$$i\hbar \hat{A}_2 = c \sum_{i=1}^d \alpha_i \left(-i\hbar \frac{\partial}{\partial r_i} \right) + mc^2 \beta \quad (18b)$$

and calculate $\hat{U}_{\hat{A}_1}(t + \tau, t, \delta)$ and $\hat{U}_{\hat{A}_2}(t + \tau, t, \delta)$ in the following way.

The operator $\hat{U}_{\hat{A}_1}(t + \tau, t, \delta)$ may be determined by splitting \hat{A}_1 further into

$$\hat{A}_1 = \hat{A}_{1,1} + \hat{A}_{1,2} \quad (19)$$

with

$$i\hbar \hat{A}_{1,1} = q\phi(\mathbf{x}, t), \quad (20a)$$

$$i\hbar \hat{A}_{1,2} = c \sum_{i=1}^d \alpha_i (-qA_i(\mathbf{x}, t)). \quad (20b)$$

Because $\hat{A}_{1,1}$ and $\hat{A}_{1,2}$ commute we may factorize the operator $\hat{U}_{\hat{A}_1}(t + \tau, t, \delta)$ into

$$\hat{U}_{\hat{A}_1}(t + \tau, t, \delta) = \hat{U}_{\hat{A}_{1,1}}(t + \tau, t, \delta) \hat{U}_{\hat{A}_{1,2}}(t + \tau, t, \delta) \quad (21)$$

with the diagonal operator

$$\hat{U}_{\hat{A}_{1,1}}(t + \tau, t, \delta) \Psi(\mathbf{x}, t) = \exp\left(-\delta \frac{i}{\hbar} \int_t^{t+\tau} q\phi(\mathbf{x}, t') dt'\right) \Psi(\mathbf{x}, t). \quad (22)$$

The operator $\hat{U}_{\hat{A}_{1,2}}(t + \tau, t, \delta)$, however, involves matrix exponentials, which may be calculated by taking into account the Dirac algebra (15). Exponentials of α_i are obtained by summing the exponential-function's Taylor sum explicitly. Introducing some auxiliary complex numbers a_i we find

$$\begin{aligned} \exp\left(i \sum_{i=1}^d a_i \alpha_i\right) &= \sum_{k=0}^{\infty} \frac{1}{k!} \left(i \sum_{i=1}^d a_i \alpha_i\right)^k \\ &= \sum_{k=0}^{\infty} \frac{(-1)^k}{(2k)!} \left(\sum_{i=1}^d a_i \alpha_i\right)^{2k} + i \sum_{i=1}^d a_i \alpha_i \sum_{k=0}^{\infty} \frac{(-1)^k}{(2k+1)!} \left(\sum_{i=1}^d a_i \alpha_i\right)^{2k} \\ &= \sum_{k=0}^{\infty} \frac{(-1)^k}{(2k)!} \left(\sqrt{\sum_{i=1}^d a_i^2}\right)^{2k} + i \sum_{i=1}^d a_i \alpha_i \sum_{k=0}^{\infty} \frac{(-1)^k}{(2k+1)!} \left(\sqrt{\sum_{i=1}^d a_i^2}\right)^{2k} \\ &= \cos\left(\sqrt{\sum_{i=1}^d a_i^2}\right) + i \sum_{i=1}^d a_i \alpha_i \left(\sum_{i=1}^d a_i^2\right)^{-1/2} \sin\left(\sqrt{\sum_{i=1}^d a_i^2}\right), \end{aligned} \quad (23)$$

where we have used

$$\begin{aligned} \left(\sum_{i=1}^d a_i \alpha_i\right)^{2k} &= \left(\left(\sum_{i=1}^d a_i \alpha_i\right)^2\right)^k \\ &= \left(\sum_{i=1}^d a_i^2 \alpha_i^2 + \sum_{i=1}^d \sum_{j=1}^{i-1} a_i a_j (\alpha_i \alpha_j + \alpha_j \alpha_i)\right)^k = \left(\sum_{i=1}^d a_i^2\right)^k. \end{aligned} \quad (24)$$

For convenience, let us define

$$\bar{A}_i(\mathbf{x}, t) = \int_t^{t+\tau} A_i(\mathbf{x}, t') dt' \quad (25a)$$

and

$$\bar{A}(\mathbf{x}, t) = \sqrt{\sum_{i=1}^d \bar{A}_i(\mathbf{x}, t)^2}, \quad (25b)$$

then we get with (23)

$$\begin{aligned} \hat{U}_{\hat{A}_{1,2}}(t + \tau, t, \delta) \Psi(\mathbf{x}, t) &= \\ \left(\cos\left(-\frac{\delta c}{\hbar} \bar{A}(\mathbf{x}, t)\right) + i \sum_{i=1}^d \frac{\bar{A}_i(\mathbf{x}, t)}{\bar{A}(\mathbf{x}, t)} \alpha_i \sin\left(-\frac{\delta c}{\hbar} \bar{A}(\mathbf{x}, t)\right)\right) \Psi(\mathbf{x}, t). \end{aligned} \quad (26)$$

The operator $\hat{U}_{\hat{A}_2}(t_2, t_1, \delta)$ equals the time evolution operator of the free particle Dirac Hamiltonian. In Fourier space it has the form

$$\hat{U}_{\hat{A}_2}(t + \tau, t, \delta) = \exp\left(-\delta \tau \frac{i}{\hbar} \left(c \sum_{i=1}^d \alpha_i p_i + mc^2 \beta\right)\right), \quad (27)$$

where p_i denotes the i th component of the momentum vector \mathbf{p} . In order to calculate the operator exponential in (27) we have to diagonalize the operator

$$i\hbar \hat{A}_2 = c \sum_{i=1}^d \alpha_i p_i + mc^2 \beta \quad (28)$$

by introducing the scalars

$$d_{\pm}(\mathbf{p}) = \left(\frac{1}{2} \pm \frac{1}{2\sqrt{1 + \mathbf{p}^2/(mc)^2}}\right)^{1/2}, \quad (29)$$

the unitary matrix

$$\hat{u}(\mathbf{p}) = d_+(\mathbf{p}) + d_-(\mathbf{p}) \sum_{i=1}^d \frac{p_i}{|\mathbf{p}|} \beta \cdot \alpha_i. \quad (30)$$

and its Hermitian adjoint $u^\dagger(\mathbf{p})$. The matrix $\hat{u}(\mathbf{p}) i\hbar \hat{A}_2 u^\dagger(\mathbf{p})$ is diagonal [21] and it reads for two- or three-dimensional systems

$$\hat{u}(\mathbf{p}) i\hbar \hat{A}_2 u^\dagger(\mathbf{p}) = \begin{pmatrix} E(\mathbf{p}) & 0 & 0 & 0 \\ 0 & E(\mathbf{p}) & 0 & 0 \\ 0 & 0 & -E(\mathbf{p}) & 0 \\ 0 & 0 & 0 & -E(\mathbf{p}) \end{pmatrix}, \quad (31)$$

with

$$E(\mathbf{p}) = \sqrt{m^2 c^4 + \mathbf{p}^2 c^2}. \quad (32)$$

Thus, the Fourier space operator (27) simplifies to

$$\begin{aligned} \hat{U}_{\hat{A}_2}(t + \tau, t, \delta) \tilde{\Psi}(\mathbf{p}, t) &= \\ \hat{u}(\mathbf{p})^\dagger \begin{pmatrix} e^{-iE(\mathbf{p})\tau/\hbar} & 0 & 0 & 0 \\ 0 & e^{-iE(\mathbf{p})\tau/\hbar} & 0 & 0 \\ 0 & 0 & e^{iE(\mathbf{p})\tau/\hbar} & 0 \\ 0 & 0 & 0 & e^{iE(\mathbf{p})\tau/\hbar} \end{pmatrix} \hat{u}(\mathbf{p}) \tilde{\Psi}(\mathbf{p}, t). \end{aligned} \quad (33)$$

For the one-dimensional Dirac equation where the vector potential reduces to a scalar $A_1(x, t)$ and the wave function $\Psi(x, t)$ has only two components the operators $\hat{U}_{\hat{A}_{1,2}}(t + \tau, t, \delta)$ and $\hat{U}_{\hat{A}_2}(t + \tau, t, \delta)$ have to form

$$\begin{aligned} \hat{U}_{\hat{A}_{1,2}}(t + \tau, t, \delta) \Psi(x, t) &= \\ \left(\cos\left(-\frac{\delta c}{\hbar} \bar{A}_1(x, t)\right) + i \alpha_1 \sin\left(-\frac{\delta c}{\hbar} \bar{A}_1(x, t)\right)\right) \Psi(x, t) \end{aligned} \quad (34)$$

and

$$\begin{aligned} \hat{U}_{\hat{A}_2}(t + \tau, t, \delta) \tilde{\Psi}(p, t) &= \\ \hat{u}(p)^\dagger \begin{pmatrix} e^{-iE(p)\tau/\hbar} & 0 \\ 0 & e^{iE(p)\tau/\hbar} \end{pmatrix} \hat{u}(p) \tilde{\Psi}(p, t). \end{aligned} \quad (35)$$

III. GPU implementations of the Fourier split operator method

III.1. GPU computing

As we have outlined in Section II, the Fourier split operator method consists of three elementary steps, the application of the two operators $\hat{U}_{\hat{A}_1}$ and $\hat{U}_{\hat{A}_2}$ plus the calculation of the fast Fourier transform and its inverse. All these three steps may be parallelized. The application of $\hat{U}_{\hat{A}_1}$ or $\hat{U}_{\hat{A}_2}$ is embarrassingly parallel—each grid point can be updated independently from others—and also the fast Fourier transform may be parallelized efficiently [22].

Scalable parallel performance, however, is difficult to attain on traditional CPU systems due to the imbalance of high CPU speed and relatively slow memory. The Fourier split operator method in particular is inherently memory bandwidth bounded because the application of $\hat{U}_{\hat{A}_1}$ or $\hat{U}_{\hat{A}_2}$ requires only $O(N)$ operations and the calculation of the Fourier transform takes $O(N \log N)$, where N denotes the number of grid points. Its low computational complexity is a very beneficial feature of the Fourier split operator method but it also means that there is no or little potential data reuse making the memory bandwidth a limiting factor.

GPU computing [2, 3] may help to overcome these limitations. Due to its massively parallel architecture, GPUs reach a peak performance in floating point number operations that is more than one order of magnitude higher than the peak performance of current CPUs. GPUs also provide higher maximal memory bandwidth. Intel's Core™ 7 CPUs, for example, reach up to 25.6 GB/s memory bandwidth while NVIDIA's Tesla™ M2050 and M2070 computing processors—which are based on GPU technology—offer up to 148 GB/s memory bandwidth, see also Table 1. Modern GPUs attain their impressive computational performance figures without significantly exceeding the power consumption of high-end CPUs. In fact, GPUs and other accelerators provide the best floating point performance per watt of current high performance computing architectures.

III.2. CUDA parallel architecture and programming model

GPU computing is enabled by programming models that provide a set of abstractions that enable to express data parallelism and task parallelism. These programming models are typically implemented by equipping a sequential general purpose programming language, as for example C or Fortran, with extensions for parallel programming and providing an application programming interface. OpenCL [23], Microsoft DirectCompute and CUDA by NVIDIA are the most popular programming models for GPU computing. For our implementation of the Fourier split operator method we chose the CUDA programming model (version 3.2) which works only

TABLE 1: Comparison of some technical key features of the high-end consumer graphics card NVIDIA GeForce GTX 480 and the computing processor module Tesla M2050 (source: NVIDIA).

	GeForce GTX 480	Tesla M2050
Processor cores	480	448
Processor core clock	1.40 GHz	1.15 GHz
Memory	1.5 GB	3 GB
Memory clock	1.848 GHz	1.546 GHz
Memory bandwidth	177 GB/s	148 GB/s
Power consumption	≈ 250 W	≤ 225 W

for graphics hardware by NVIDIA but provides more flexibility than OpenCL or DirectCompute. OpenCL is the platform independent but has no support for C++ and DirectCompute works only on Microsoft Windows systems.

The term CUDA also refers to some GPU architectures by the hardware manufacturer NVIDIA. The latest CUDA GPU architecture is called Fermi. It features up to 512 computing cores which are organized in 16 streaming multiprocessors of 32 cores each. Each of the 32 cores of a streaming multiprocessor executes the same instruction on different data sets at the same time. Therefore, GPUs belong to the class *Single Instruction, Multiple Data streams* (SIMD) in Flynn's taxonomy [24].

GPU accelerator cards come in two different flavors, traditional graphics cards, as the NVIDIA GeForce GTX 480, and dedicated computing processors, as the NVIDIA Tesla M2050. Table 1 compares some technical key features of these two cards. Both are based on the so-called Fermi architecture but differ, for example, in memory and clock rates. According to this table, the dedicated computing processor M2050 seems to be inferior to the GTX 480. However there are two important features that qualify the M2050 for high performance computing applications. It has larger memory with error correction and full double precision performance. In consumer Fermi GPUs as the GTX 480, the number of double precision floating point operations that may be carried out per clock cycle is reduced by a factor of four as compared to dedicated GPU computing processors based on the Fermi architecture.

A GPU has its own memory. This means, before one can run a GPU computation the input data has to be transferred from the computer's main memory (also called host memory) to the GPU device memory. GPU computations are carried out in device memory and final results will be transferred back into host memory. Memory transfer is a relatively slow operation. Therefore, one should reduce the number of memory transfers to reach high performance. The CUDA architecture provides three different kinds of host memory, non-pinned, pinned, and mapped memory. Memory transfers with pinned memory are faster than memory transfers with non-pinned memory. However, allocating pinned memory is slow. Mapped memory allows a GPU to work directly on data in host memory. Thus, with mapped memory there is no need to copy data but a GPU will access data in mapped memory not as fast as in device memory. We will study the impact of memory transfer in section IV.1.

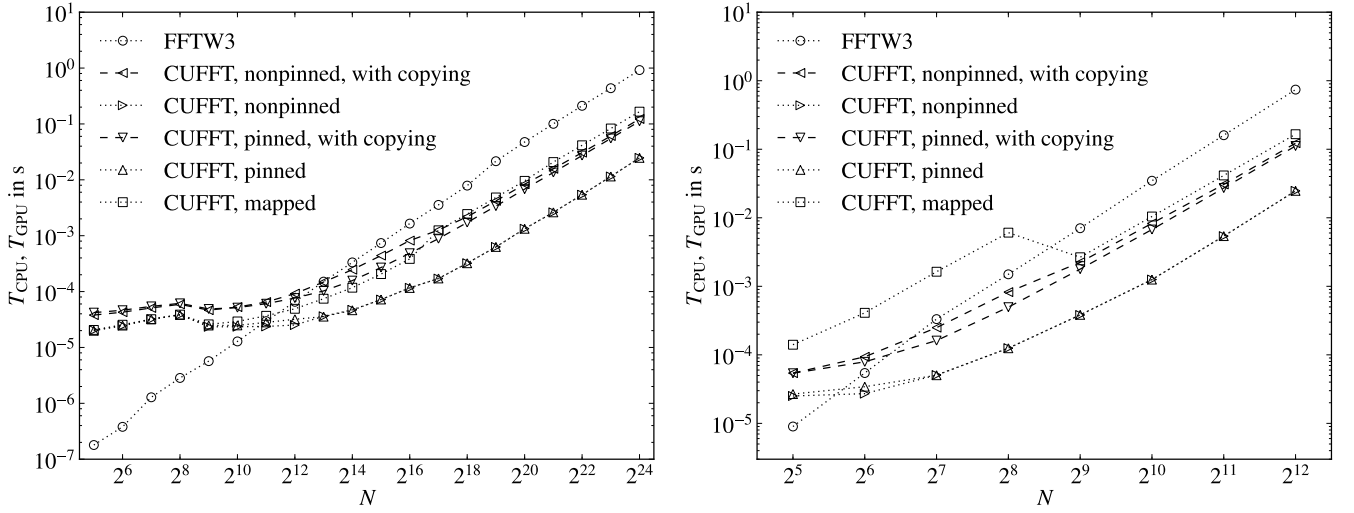


FIG. 1: Wall clock time to compute a one-dimensional (left panel) and two-dimensional (right panel) fast Fourier transform of complex-valued double precision arrays of size N and $N \times N$, respectively. Computations are done in place, that is, the input data is overwritten by the output. Times T_{CPU} and T_{GPU} denote time for computation utilizing a single CPU core (Intel Core i7 CPU at 2.93 GHz) and the FFTW3 library and a GeForce GTX 480 GPU and the CUFFT library, respectively. In the GPU case, measurements were carried out including and excluding data transfer between host memory (nonpinned or pinned) and GPU memory. Mapped memory resides on the host eliminating the need for data transfer.

III.3. Implementation

We have developed highly tuned codes for the Fourier split operator method that allow us to propagate Schrödinger wave functions and Dirac wave functions in one and two dimensions. For each equation and each dimension we implemented a conventional non-parallel CPU code as well as a CUDA code that performs all computations on a GPU. To carry out the fast Fourier transform CPU codes utilize the FFTW3 library [25]. The GPU codes employ the CUFFT library for fast Fourier transforms and comprise light-weight kernels that accomplish the action of the operators \hat{U}_{A_1} and \hat{U}_{A_2} for each grid point in parallel.

IV. Performance results

IV.1. Fast Fourier transform

The overall performance of the Fourier split operator method is mainly determined by the performance of the fast Fourier transform. Thus, we compare in Figure 1 the wall clock time T_{CPU} that is required to calculate the fast Fourier transform on a CPU with the time T_{GPU} to perform the same task on a GPU. For not too small problems the GPU outperforms the CPU significantly. With our hardware setup (Intel Core i7 CPU at 2.93 GHz with a GeForce GTX 480 GPU) we got a speedup up to a factor of about 30. For small problems, however, the overhead of starting and coordinating parallel CUDA threads prevents to archive good fast Fourier transform performance on GPUs. There is a critical problem size where the GPU

starts to outperform a CPU in calculating a fast Fourier transform, which is for our hardware setup at data sets of about 2^{12} complex numbers.

For GPU computations, one may use nonpinned, pinned, or mapped memory on the host. If one includes the time to copy data from host memory to device memory and back again after the calculation of the Fourier transform in the measurements then the time depends on the kind of memory. Pinned memory is faster than nonpinned memory. Mapped memory eliminates the need to copy data between host and device, however, performing the fast Fourier transform on device memory plus copying is usually faster than calculating it in mapped host memory without data transfer. Only for one-dimensional Fourier transforms of intermediate size carrying out the fast Fourier transform in mapped host memory requires less time than performing the same task on device memory plus data transfer or using a (nonparallel) CPU implementation.

We took also fast Fourier transform performance measurements for a Tesla M2050 computing processor. Despite the fact that the M2050 can perform four times more double precision operations per clock cycle than the consumer GPU GeForce GTX 480, the Fourier transform performance of the Tesla M2050 is slightly lower than for the GeForce GTX 480. For a two-dimensional Fourier transform of a 4096×4096 array, for example, the Tesla M2050 reached only about 95 % of the GeForce GTX 480 GPU performance. For smaller arrays the performance difference was even larger. This may be interpreted as that the fast Fourier transform performance is bound by the GPU's memory bandwidth rather than by its floating point performance.

Our findings for the fast Fourier transform performance are consistent with the results presented in the literature. Reference [26] reports a performance gain up to a factor of about

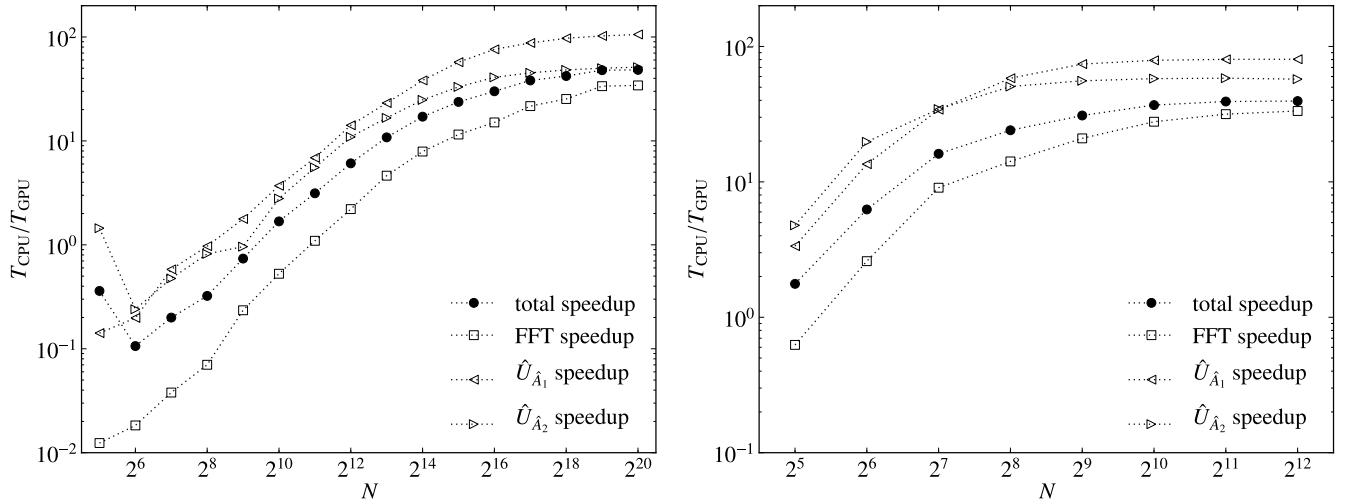


FIG. 2: Speedup $T_{\text{CPU}}/T_{\text{GPU}}$ for propagating a wave function over 128 time steps as a function of the grid size. The left panel depicts results for the one-dimensional Schrödinger equation of a grid of N data points while the right panel shows results for the two-dimensional Dirac equation of a grid of $N \times N$ data points. Performance measurements carried out on the same hardware as in Figure 1 and in double precision.

11 in calculating two-dimensional single precision fast Fourier transforms on a NVIDIA GeForce 9800 GX2 GPU[27] instead on a conventional CPU. In [22] a 8- to 40-fold improvement was archived over highly tuned CPU routines by using an algorithm that efficiently exploits GPU shared memory.

To sum up, the CUFFT GPU implementation of the fast Fourier transform is capable to outperform traditional CPU implementations by more than one order of magnitude. Data transfer between host memory and device memory poses a non-negligible overhead and should be avoided if possible or reduced by using pinned memory.

IV.2. Fourier split operator method

In order to determine the speedup that may be attained by switching from a CPU implementation to a GPU implementation of the Fourier split operator method, we propagated one- and two-dimensional Gaussian wave packets in a harmonic potential over 128 time steps under the effect of the Schrödinger equation (9) and the Dirac equation (14), receptively. We measured

- the overall time to propagate 128 time steps (In the case of GPU computations, this includes data transfer between host memory and device memory before the first step and after the last step.),
- the time to perform the fast Fourier transform plus its inverse,
- the time to apply the operator $\hat{U}_{\hat{A}_1}$ in position space, and
- the time to apply the operator $\hat{U}_{\hat{A}_2}$ in Fourier space

for various grids of size N and $N \times N$. Figure 2 shows the speedup $T_{\text{CPU}}/T_{\text{GPU}}$ for these four different tasks for the one-dimensional Schrödinger equation and for the two-dimensional Dirac equation as a function of the grid size. If

one considers the individual steps of the Fourier split operator method, then one finds that the speedup for the application of $\hat{U}_{\hat{A}_1}$ and $\hat{U}_{\hat{A}_2}$ reaches between 50 to about 100 for large systems (more than about 2^{18} grid points). For the fast Fourier transform, however, we get a speedup of only about 30 limiting the overall speedup to about 40 for large systems, which is still a significant improvement over the CPU implementation. Results for the two-dimensional Schrödinger equation and the one-dimensional Dirac equation are not shown in Figure 2 because they are qualitatively very similar to the two-dimensional Dirac equation and the one-dimensional Schrödinger equation, respectively.

V. Applications

In this section we will show some applications of our Fourier split operator GPU codes for the solution of the time dependent Dirac equation. With conventional CPU codes [16], these kinds of applications would require more than an order of magnitude more computing time or may not be performed in an admissible amount of time. For our numerical simulation we adopt the atomic unit system (a. u.) that is established on the Bohr radius, the electron's mass, the reduced Planck constant and the absolute value of the electron's charge as its base units.

V.1. Evolution of a free Gaussian wave packet

A d -dimensional wave packet may be formed by superimposing plane waves solutions of the Dirac equation with defined

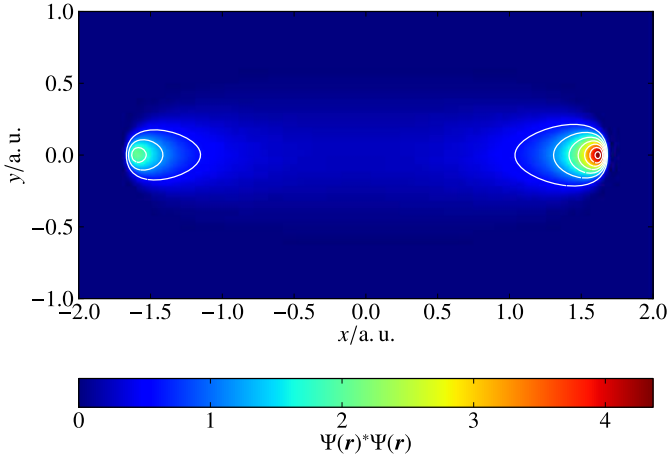


FIG. 3: Probability density of a Dirac wave packet with narrow asymmetric momentum distribution at time $t = 0.0125$ a.u., see text for specific parameters. Due to the finite speed of light, the wave packet splits into two shock fronts traveling into opposite directions.

momentum \mathbf{p} , viz.

$$\Psi_{\text{packed}}(\mathbf{x}) = \frac{1}{(2\pi\hbar)^{d/2}} \int \rho(\mathbf{p}) \exp\left(\frac{i\mathbf{x} \cdot \mathbf{p}}{\hbar}\right) d^d p \quad (36)$$

The quantity $\rho(\mathbf{p})$ determines the momentum distribution. For a two-dimensional Gaussian distribution with mean momentum $\bar{\mathbf{p}}$ and momentum widths σ_f in forwards direction and σ_s in sideways direction the function $\rho(\mathbf{p})$ reads

$$\rho_{\text{Gaussian}}(\mathbf{p}) = u(\mathbf{p}) \frac{1}{(2\pi)^{1/2} |\Sigma|^{1/4}} \exp\left(-\frac{(\mathbf{p} - \bar{\mathbf{p}})^T \Sigma^{-1} (\mathbf{p} - \bar{\mathbf{p}})}{4}\right), \quad (37)$$

where $u(\mathbf{p})$ denotes a column of the matrix $\tilde{u}(\mathbf{p})$ (30) (selecting one of the four momentum eigenstates with momentum \mathbf{p}) and with the matrix

$$\Sigma^{-1} = R \begin{pmatrix} 1/\sigma_f^2 & 0 \\ 0 & 1/\sigma_s^2 \end{pmatrix} R^{-1} \quad (38)$$

and

$$R = \begin{pmatrix} \cos(\phi) & -\sin(\phi) \\ \sin(\phi) & \cos(\phi) \end{pmatrix} \quad \text{and} \quad \tan \phi = \frac{\bar{p}_y}{\bar{p}_x}. \quad (39)$$

The evolution of the probability density of a relativistic Dirac wave packet differs significantly from the non-relativistic theory of the Schrödinger equation, where a Gaussian wave packet broadens but remains its Gaussian shape for all times and the maximum of the probability density travels with the same speed as the motion of the center of mass does. We consider the propagation of a relativistic Gaussian Dirac wave packet in two dimensions with asymmetric momentum widths $\sigma_f = 200$ a.u. and $\sigma_s = 20$ a.u., a mean momentum $\bar{\mathbf{p}} = (40 \text{ a.u.}, 0 \text{ a.u.})$, positive energy and spin up. This corresponds to a rather narrow initial wave packet. The probability density of this two-dimensional relativistic

Dirac wave packet at time $t = 0.0125$ a.u. is illustrated in FIG. 3. The position of the initial probability density's maximum corresponds to the initial center of mass at the center of the coordinate system. The wave packet's center of mass moves in accordance with classical predictions with velocity $|\bar{\mathbf{p}}|/(m\sqrt{1 + \bar{\mathbf{p}}^2/(mc)^2}) \approx 38$ a.u. along the x -axis. Due to the finite speed of light [28], however, shock fronts emerge traveling approximately with the speed of light into the forward and backward directions. The maxima of the probability density no longer coincide with the center of mass. See [29, 30] for an investigation of similar relativistic effects.

V.2. Eigenstates

Propagating a trial wave function $\Psi(\mathbf{x}, t)$ in a time independent potential from $t = 0$ to some later time $t = T$ allows us to determine the potential's eigenenergies [12] by calculating the autocorrelation function

$$\chi(t) = \int \Psi(\mathbf{x}, 0)^* \Psi(\mathbf{x}, t) d^d x. \quad (40)$$

The Fourier transform of the autocorrelation function

$$\tilde{\chi}(E) = \int_0^T \chi(t) (1 - \cos(2\pi t/T)) \exp(iEt/\hbar) dt \quad (41)$$

exhibits pronounced peaks at the eigenenergies provided that the initial trial function is not orthogonal to some eigenfunction. Once we have determined an eigenenergy E from the spectrum $\tilde{\chi}(E)$, the potential's eigenfunction $\Psi_E(\mathbf{x})$ (assuming there is no degeneracy) with eigenenergy E results from the integral

$$\Psi_E(\mathbf{x}) \sim \int_0^T \Psi(\mathbf{x}, t) (1 - \cos(2\pi t/T)) \exp(iEt/\hbar) dt. \quad (42)$$

For one- or two-dimensional model systems it is common practice to mimic the Coulomb potential

$$V_C(\mathbf{x}) = -e\phi(\mathbf{x}) = -\frac{Ze^2}{4\pi\epsilon_0|\mathbf{x}|} \quad (43)$$

by a soft-core potential

$$V_{\text{sc}}(\mathbf{x}) = -e\phi(\mathbf{x}) = -\frac{Ze^2}{4\pi\epsilon_0 \sqrt{\mathbf{x}^2 + (\zeta/Z)^2}}. \quad (44)$$

In (43) and (44) e denotes the elementary charge, Z the atomic number, ϵ_0 the vacuum permittivity, and ζ the soft-core parameter.

If one replaces the three-dimensional Coulomb potential (43) by a lower dimensional soft-core potential (44), it is often appropriate to choose the soft-core parameter ζ such that both potential share the same groundstate energy. In Schrödinger theory, the Coulomb potential (43) has a groundstate energy of $-mc^2(\alpha Z)^2/2$, where we have introduced the electron mass m and the fine structure constant $\alpha = e^2/(4\pi\epsilon_0\hbar c)$. Numerically

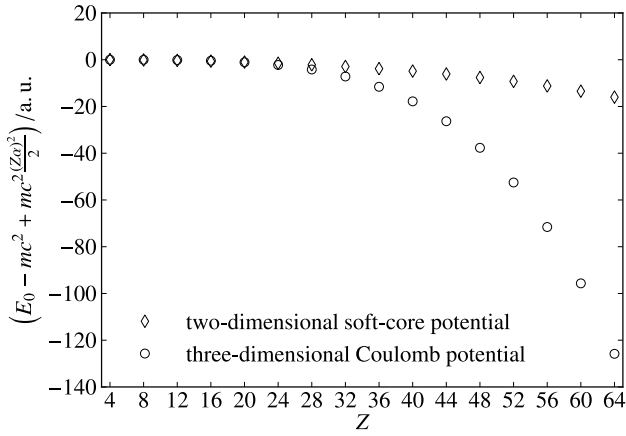


FIG. 4: Relativistic corrections to the ground state energy E_0 for the two-dimensional soft-core potential (44) with $\zeta = 0.791$ and the three-dimensional soft-core potential for different values of the atomic number Z . In the non-relativistic Schrödinger theory, both potentials share approximately the same ground state energy of $-mc^2(\alpha Z)^2/2$.

we find that for $\zeta \approx 1.4133$ the one-dimensional soft-core potential has approximately the same groundstate energy as the three-dimensional Coulomb potential with the same atomic number. In two dimensions, one has to set $\zeta \approx 0.791$ in order to match the groundstate energies. These statements hold for all Z .

In Dirac's quantum theory, the groundstate energy of the three-dimensional Coulomb potential (43) is $mc^2 \sqrt{1 - (\alpha Z)^2}$ and the soft-core parameter that matches the groundstate energies depends on Z , however, it is close to the values for the Schrödinger case. The Dirac groundstate energy of the Coulomb potential equals the Schrödinger groundstate energy plus the rest mass energy mc^2 and relativistic corrections proportional in Z^4 , viz.

$$mc^2 \sqrt{1 - (\alpha Z)^2} = mc^2 - mc^2 \frac{(\alpha Z)^2}{2} - mc^2 \frac{(\alpha Z)^4}{8} + \dots \quad (45)$$

We determined the Dirac groundstate energy E_0 of the two-dimensional soft-core potential (44) with $\zeta = 0.791$ as a function of the atomic number Z with high numerical accuracy and found that relativistic corrections to the two-dimensional soft-core-potential's groundstate energy are weaker than for the three-dimensional Coulomb potential as shown in Fig. 4.

V.3. Free wave packed scattering at a light pulse

At high laser intensities, charged particles in electromagnetic waves are not only affected by the electric field component but also by the magnetic field component. At intensities larger than about $55m\hbar\omega^2/(c\mu_0q^2)$, the Lorentz force becomes relevant and at intensities above $0.1m^2c^2\omega^2/q^2$ also relativistic effects have to be taken into account [31], here ω denotes the laser's angular frequency and μ_0 the vacuum permeability.

We simulated the two-dimensional motion of a free wave packed in a laser pulse with the electromagnetic fields

$$\mathbf{E}(\mathbf{x}, t) = E_0 \sin(\mathbf{k} \cdot \mathbf{x} - \omega t) f_{j,l}(\mathbf{k} \cdot \mathbf{x} - \omega t), \quad (46a)$$

$$\mathbf{B}(\mathbf{x}, t) = \frac{E_0}{c} \sin(\mathbf{k} \cdot \mathbf{x} - \omega t) f_{j,l}(\mathbf{k} \cdot \mathbf{x} - \omega t) \quad (46b)$$

with an envelope function $f_{j,l}(\eta)$ of j half cycles with a linear turn-on ramp and a linear turn-off ramp of l half cycles and a constant plateau in between, viz.

$$f_{j,l}(\eta) = \begin{cases} (\eta + j\pi)/(l\pi) & \text{if } -j \leq \eta/\pi \leq -j + l, \\ 1 & \text{if } -j + l \leq \eta/\pi \leq -l, \\ -\eta/(l\pi) & \text{if } -l \leq \eta/\pi \leq 0, \\ 0 & \text{else} \end{cases} \quad (46c)$$

and $|\mathbf{k}| = 2\pi/\lambda = \omega/c$ and $\mathbf{E}_0 \perp \mathbf{k}$. Figure 5 shows the wave-function's probability density at different times for a free wave packed scattering at a strong laser pulse traveling along the x -direction and having an overall length of eight half cycles and turn-on/off ramps of two half cycles. Its field strength is $|\mathbf{E}_0| = 3000$ a.u. and its wave length $\lambda = 40$ a.u. The electric field component accelerates the wave packet back and forth along the y -direction, while the Lorentz force causes a drift into the x -direction.

V.4. Ionization

In our fourth example, we consider ionization in ultra-strong and ultra-short laser pulses. The Dirac ground state wave packet of a soft-core potential (44) with $Z = 32$ and $\zeta = 0.791$ is excited by an external laser pulse (46) with wavelength $\lambda = 10$ a.u. and peak electric field strength $|\mathbf{E}_0| = 3072$ a.u.. We consider a short laser pulse of four half-cycles including a turn-on and a turn-off half-cycle. Figure 6 shows the electron density after the laser pulse has passed the atomic core. In the setup, we have chosen here, the laser travels from left to right and the electric field points into the y -direction, accelerating the wave packet into the positive y -direction in the first and third half-cycles and into the negative y -direction during the second and fourth half-cycles. The Lorentz force causes an acceleration into the laser's propagation direction. The strong interference patterns in Fig. 6 form by the interaction of wave-function's different components traveling into different directions.

Note the huge wave function spreading during the ionization. While the ground state wave function has a width of about $1/32$ a.u., the wave function after ionization spreads over several atomic units as shown in Fig. 6. Wave function spreading poses a major challenge in the numerical simulation of light matter interaction at high intensities limiting the simulations to very short time scales where the size of the wave function remains of manageable size. The simulation of the ionization process of Fig. 6 took only about ten minutes using our GPU Fourier split operator implementation. Thus, the simulation of even longer interactions with laser pulses of longer

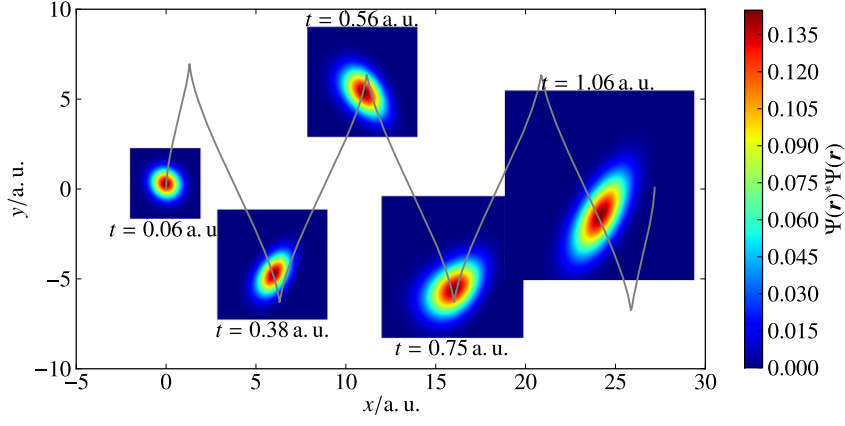


FIG. 5: Free wave packed scattering at a strong laser pulse. The false color plots show the wave-packet's probability density at different points in time t . The solid gray line indicates the center of mass trajectory. The laser pulse travels from left to right. See text for detailed parameters. Note that the computational grid follows the center of mass motion.

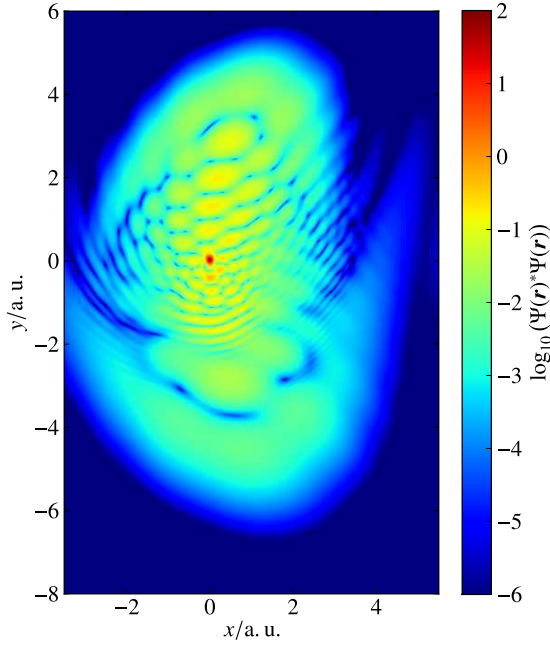


FIG. 6: A wave-function's probability density after ionization from the ground state of a soft-core potential (44) by a ultra-strong laser pulse, see text for details.

wave lengths has become feasible by our GPU implementation.

VI. Conclusions and outlook

In this contribution we evaluated GPUs as a massively parallel computing architecture for the solution of time dependent partial differential equations by means of the Fourier split operator method. The computational complexity of the Fourier split operator method is dominated by the computational complexity of the fast Fourier transform. We demonstrated that GPUs reach much better performance in computing the fast Fourier transform than current CPUs. Depending on the problem size, the performance gain may exceed one order of magnitude as compared to sequential CPU implementations. Thus, the Fourier split operator method may be implemented very efficiently on GPU architectures as we demonstrated for the time dependent Schrödinger equation and the time dependent Dirac equation. The combination of a highly parallel architecture with a high-throughput memory makes graphics processing units a very attractive architecture for implementing the Fourier split operator method. Best performance is attained if all steps of the Fourier split operator method are carried out by the GPU avoiding data transfer between host memory and GPU memory.

The fast Fourier transform is a core building block for the solution of partial differential equations as well as of many other problems from computational physics, signal processing, tomography, computational finance and other fields. Thus, we expect that also these problems may be solved much more efficiently on GPU architectures than on conventional CPUs. Taking into account that GPU computing is a rather young field it is supposed that there is still plenty potential for GPU technology and codes to mature further and to find new applications.

Acknowledgments

We would like to thank Martin Wolf and Sven Ahrens for stimulating discussions.

-
- [1] D. Kirk, W. Hwu, *Programming Massively Parallel Processors: A Hands-on Approach*, Morgan Kaufmann, 2010.
- [2] V. Kindratenko, R. Wilhelmson, R. Brunner, T. J. Martinez, W. Hwu, High-performance computing with accelerators, *Computing in Science & Engineering* 12 (4) (2010) 12–16. doi:10.1109/MCSE.2010.88.
- [3] J. Owens, M. Houston, D. Luebke, S. Green, J. Stone, J. Phillips, GPU computing, *Proceedings of the IEEE* 96 (5) (2008) 879–899. doi:10.1109/JPROC.2008.917757.
- [4] H. Sutter, A fundamental turn toward concurrency in software, *Dr. Dobbs's Journal* 30 (3) (2005) 202–210.
- [5] H. Risken, *The Fokker-Planck Equation*, 2nd Edition, Vol. 18 of Springer Series in Synergetics, Springer, 1989.
- [6] W. Harshawardhan, Q. Su, R. Grobe, Numerical solution of the time-dependent Maxwell's equations for random dielectric media, *Physical Review E* 62 (6) (2000) 8705–8712. doi:10.1103/PhysRevE.62.8705.
- [7] B. Thaller, *The Dirac Equation*, Texts and Monographs in Physics, Springer, 1992.
- [8] H. Feshbach, F. Villars, Elementary relativistic wave mechanics of spin 0 and spin 1/2 particles, *Reviews of Modern Physics* 30 (1) (1958) 24–45. doi:10.1103/RevModPhys.30.24.
- [9] P. Pechukas, J. C. Light, On the exponential form of time-displacement operators in quantum mechanics, *The Journal of Chemical Physics* 44 (10) (1966) 3897–3912. doi:10.1063/1.1726550.
- [10] H. O. Fattorini, *The Cauchy Problem*, Vol. 18 of Encyclopedia of Mathematics and its Applications, Cambridge University Press, 1985.
- [11] J. A. Fleck, J. R. Morris, M. D. Feit, Time-dependent propagation of high energy laser beams through the atmosphere, *Applied Physics* 10 (2) (1976) 129–160.
- [12] M. D. Feit, J. A. Fleck, Jr., A. Steiger, Solution of the Schrödinger equation by a spectral method, *Journal of Computational Physics* 47 (3) (1982) 412–433. doi:10.1016/0021-9991(82)90091-2.
- [13] A. D. Bandrauk, H. Shen, Exponential split operator methods for solving coupled time-dependent Schrödinger equations, *Journal of Chemical Physics* 99 (2) (1993) 1185–1193. doi:10.1063/1.465362.
- [14] J. W. Braun, Q. Su, R. Grobe, Numerical approach to solve the time-dependent Dirac equation, *Physical Review A* 59 (1) (1999) 604–612. doi:10.1103/PhysRevA.59.604.
- [15] G. R. Mocken, C. H. Keitel, Quantum dynamics of relativistic electrons, *Journal of Computational Physics* 199 (2) (2004) 558–588. doi:10.1016/j.jcp.2004.02.020.
- [16] G. R. Mocken, C. H. Keitel, FFT-split-operator code for solving the Dirac equation in 2 + 1 dimensions, *Computer Physics Communications* 178 (11) (2008) 868–882. doi:10.1016/j.cpc.2008.01.042.
- [17] S. X. Hu, C. H. Keitel, Dynamics of multiply charged ions in intense laser fields, *Physical Review A* 63 (5) (2001) 053402. doi:10.1103/PhysRevA.63.053402.
- [18] J. Javanainen, J. Ruostekoski, Symbolic calculation in development of algorithms: split-step methods for the Gross-Pitaevskii equation, *Journal of Physics A* 39 (12) (2006) L179–L184. doi:10.1088/0305-4470/39/12/L02.
- [19] G. M. Muslu, H. A. Erbay, Higher-order split-step Fourier schemes for the generalized nonlinear Schrödinger equation, *Mathematics and Computers in Simulation* 67 (6) (2005) 581–595. doi:10.1016/j.matcom.2004.08.002.
- [20] G. Strang, On the construction and comparison of difference schemes, *SIAM Journal on Numerical Analysis* 5 (3) (1968) 506–517. doi:10.1137/0705041.
- [21] B. Thaller, *Advanced Visual Quantum Mechanics*, Springer, 2004.
- [22] N. K. Govindaraju, B. Lloyd, Y. Dotsenko, B. Smith, J. Manferdelli, High performance discrete Fourier transforms on graphics processors, in: *Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, Piscataway, NJ, USA, 2008, pp. 1–12.
- [23] Khronos OpenCL Working Group, *The OpenCL Specification Version 1.1* (2010).
- [24] M. J. Flynn, Some computer organizations and their effectiveness, *IEEE Transactions on Computers* C-21 (9) (1972) 948–960. doi:10.1109/TC.1972.5009071.
- [25] M. Frigo, S. G. Johnson, The design and implementation of FFTW3, *Proceedings of the IEEE* 93 (2) (2005) 216–231. doi:10.1109/JPROC.2004.840301.
- [26] V. Surkov, Parallel option pricing with Fourier space time-stepping method on graphics processing units, *Parallel Computing* 36 (7) (2010) 372–380. doi:10.1016/j.parco.2010.02.006.
- [27] Two hardware generations older than the GeForce GTX 480 that we used.
- [28] B. Thaller, *Advanced Visual Quantum Mechanics*, Springer, 2004.
- [29] M. Ruf, H. Bauke, C. H. Keitel, A real space split operator method for the Klein-Gordon equation, *Journal of Computational Physics* 228 (24) (2009) 9092–9106. doi:10.1016/j.jcp.2009.09.012.
- [30] Q. Su, B. Smetanko, R. Grobe, Relativistic suppression of wave packet spreading, *Optics Express* 2 (7) (1998) 277–281.
- [31] H. R. Reiss, Dipole-approximation magnetic fields in strong laser beams, *Physical Review A* 63 (2000) 013409. doi:10.1103/PhysRevA.63.013409.